



US005425020A

**United States Patent** [19]

Gregg et al.

[11] Patent Number: **5,425,020**[45] Date of Patent: **Jun. 13, 1995****[54] SKEW MEASUREMENT FOR RECEIVING FRAME-GROUPS**

[75] Inventors: Thomas A. Gregg, Highland; Joseph M. Hoke, Millerton; Patrick J. Sugrue, Hyde Park, all of N.Y.

[73] Assignee: International Business Machines Corporation, Armonk, N.Y.

[21] Appl. No.: 147,701

[22] Filed: Nov. 4, 1993

[51] Int. Cl.<sup>6</sup> ..... H04J 3/14

[52] U.S. Cl. .... 370/17; 370/91; 370/100.1; 370/105

[58] Field of Search ..... 370/17, 91, 92, 105, 370/105.3, 100.1

**[56] References Cited****U.S. PATENT DOCUMENTS**

4,809,259 2/1989 Jönsson ..... 370/100.1  
 5,003,558 3/1991 Gregg et al. .... 375/108  
 5,025,458 6/1991 Casper et al. .... 375/114  
 5,267,240 11/1993 Bartow et al. .... 371/1

Primary Examiner—Douglas W. Olms

Assistant Examiner—Russell W. Blum

Attorney, Agent, or Firm—Lynn L. Augspurger

**[57] ABSTRACT**

A system and method for asynchronously transmitting data blocks, in parallel, across multiple fibers in a serial manner. Frame groups are provided as a mechanism to transmit associated data serially on each fiber and tie the data being transmitted together. The frame groups do not have sequence numbers, therefore, the receiver determines which frames are part of a frame group by the arrival times of the individual frames. The transceivers for each member of the parallel bus asynchronously achieve synchronism from either end of the fiber. A receiver of a transceiver determines a relative skew on a conductor of a link by performing skew measurements on a series of calibration frames sent by a transmitter on another end of said link. A skew monitor performs skew measurements and includes a plurality of skew measurement registers for measuring the relative skew between frames of a frame group, each register having an input from a skew counter, and outputs coupled to a microprocessor controlling the monitor. Each register gates in the value in the skew counter at the time a link control word is received. The receiver can determine that one or more frames of the calibration frames have been lost and have not exceeded the maximum measureable skew limits.

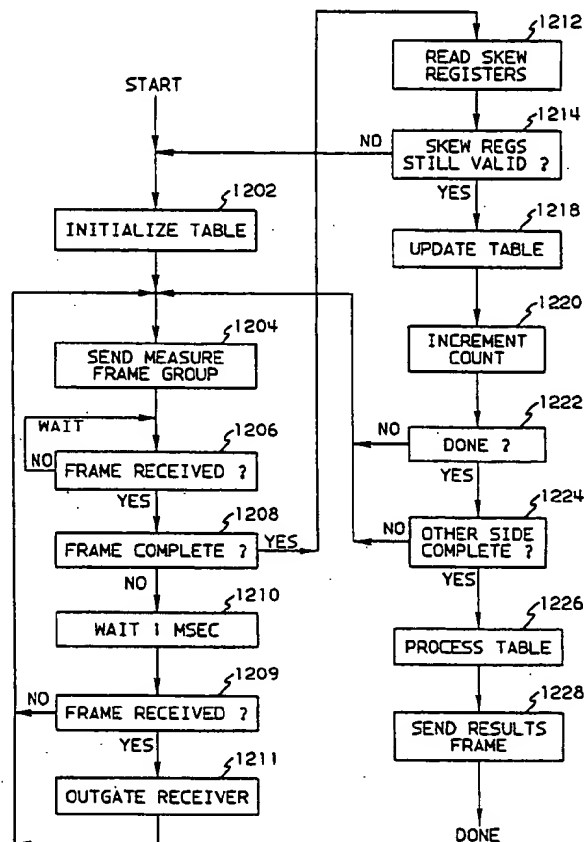
**10 Claims, 11 Drawing Sheets**

FIG. 1




FIG. 2




FIG. 3

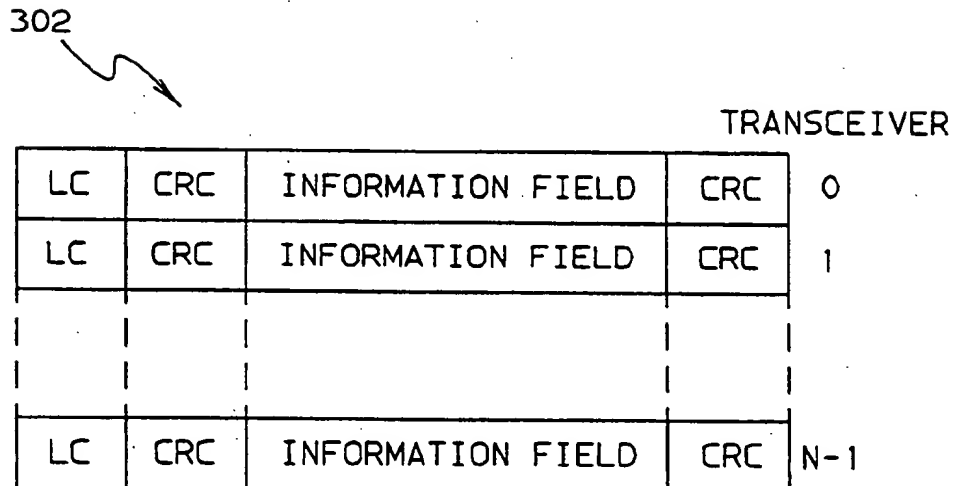
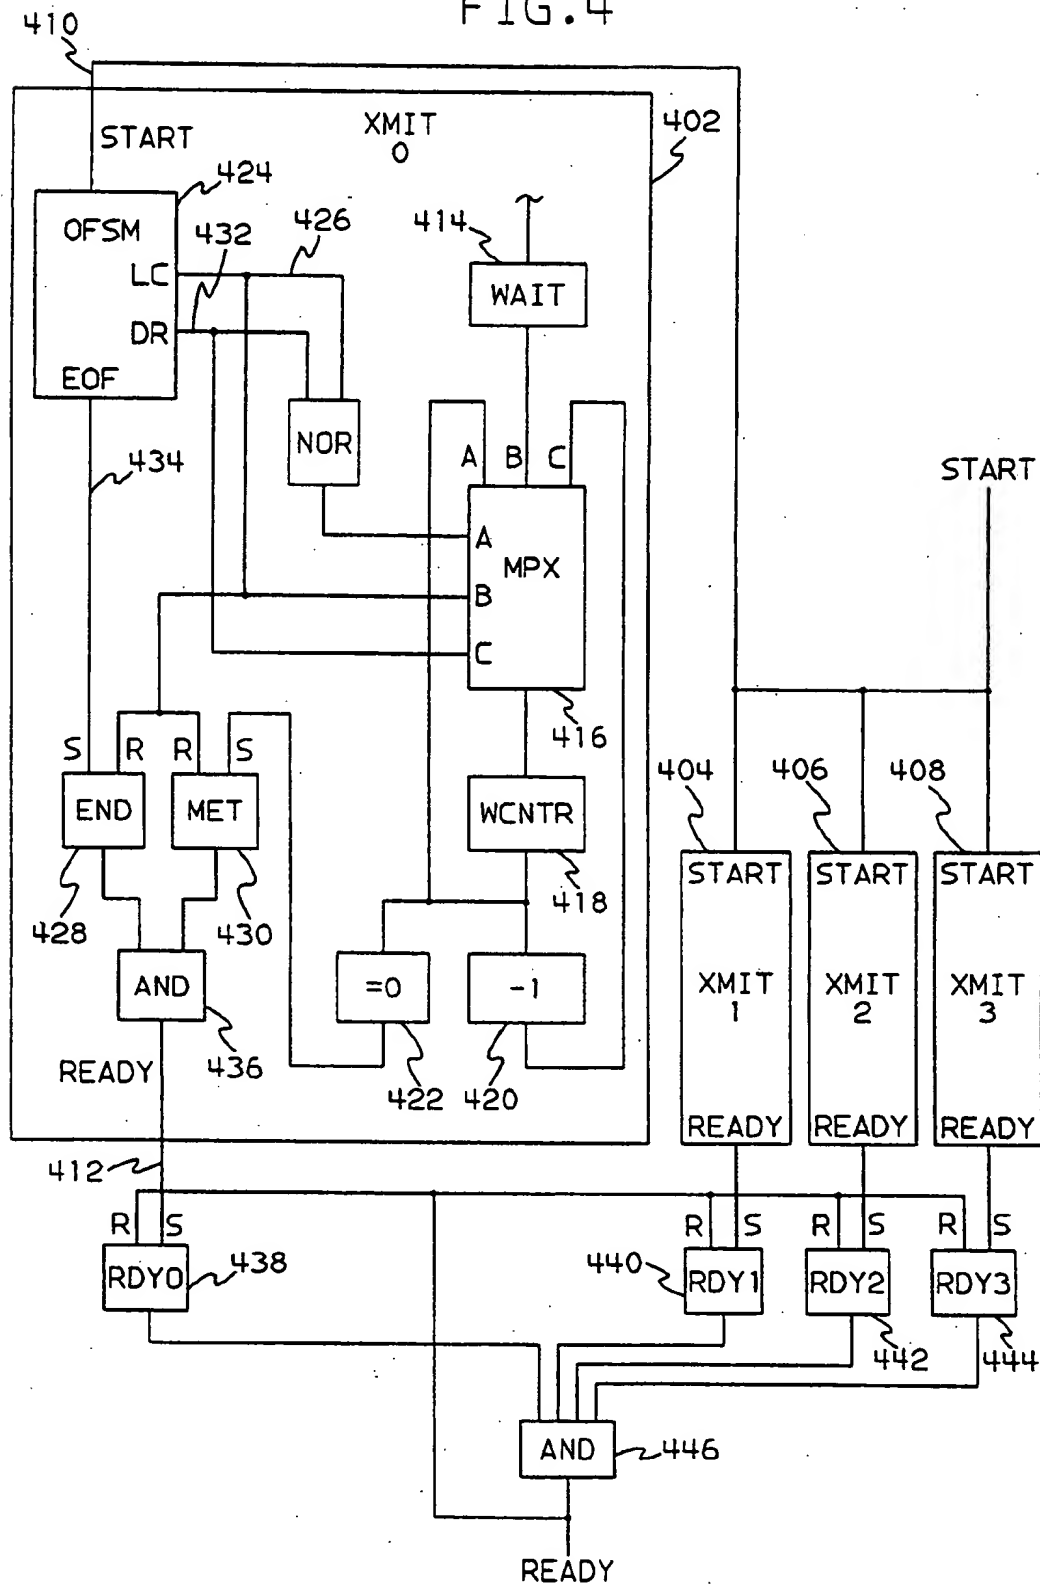
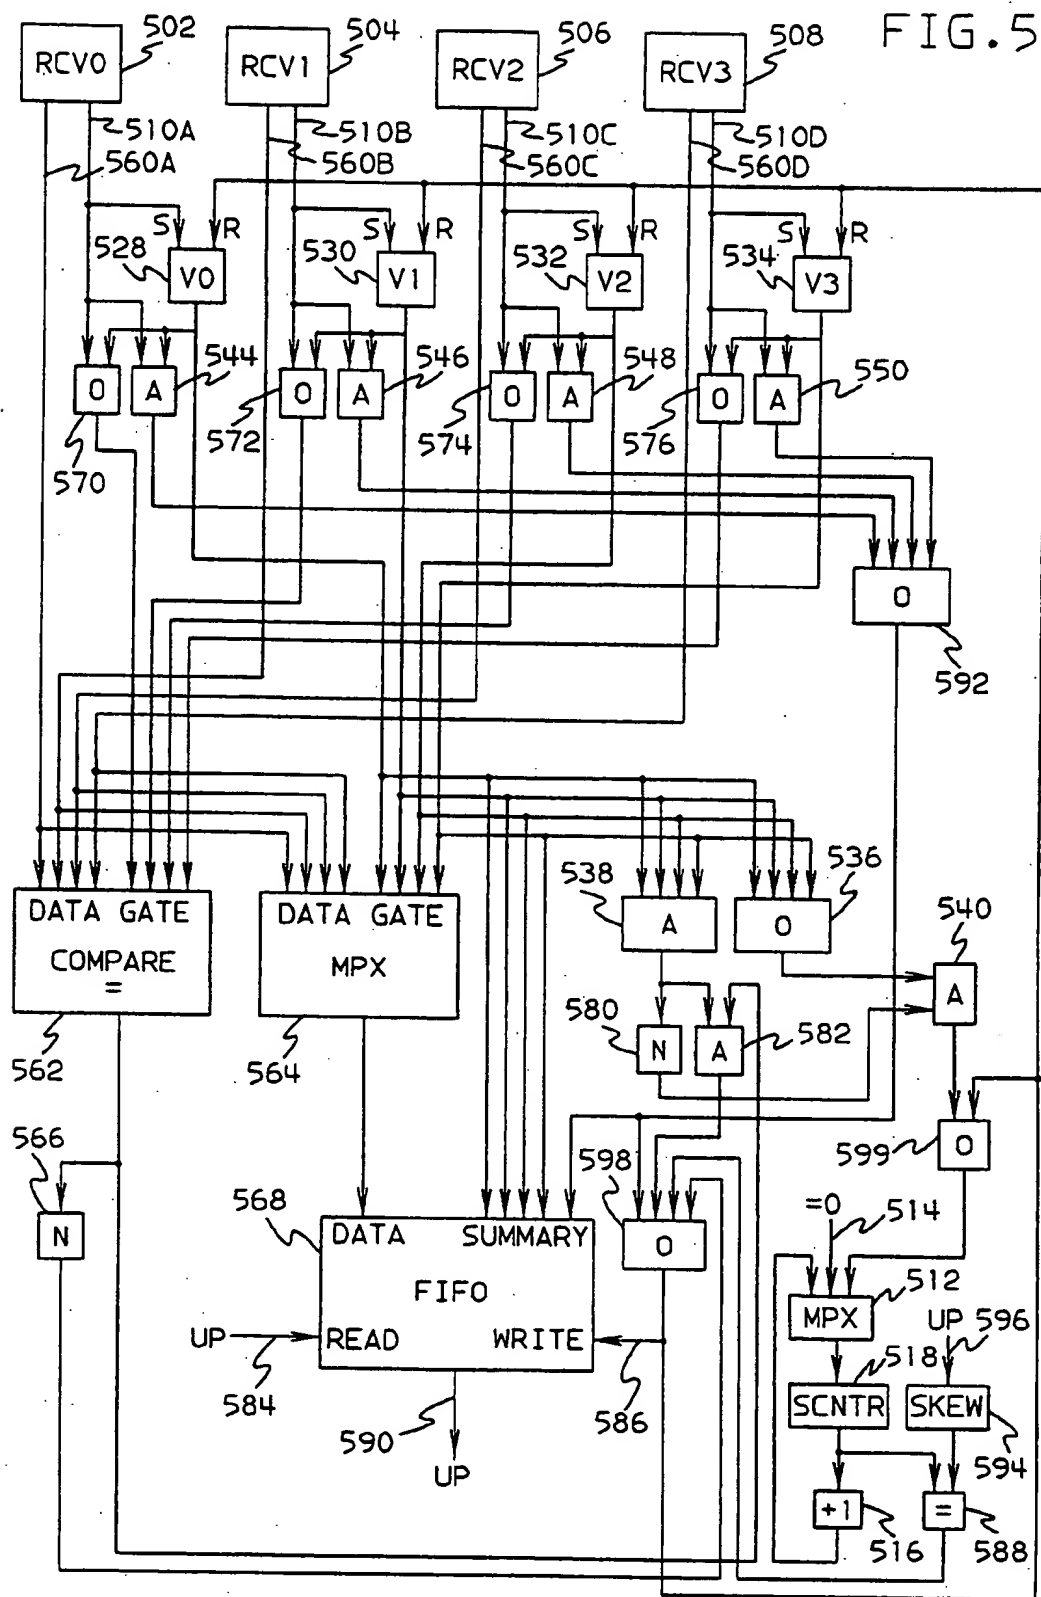


FIG. 4





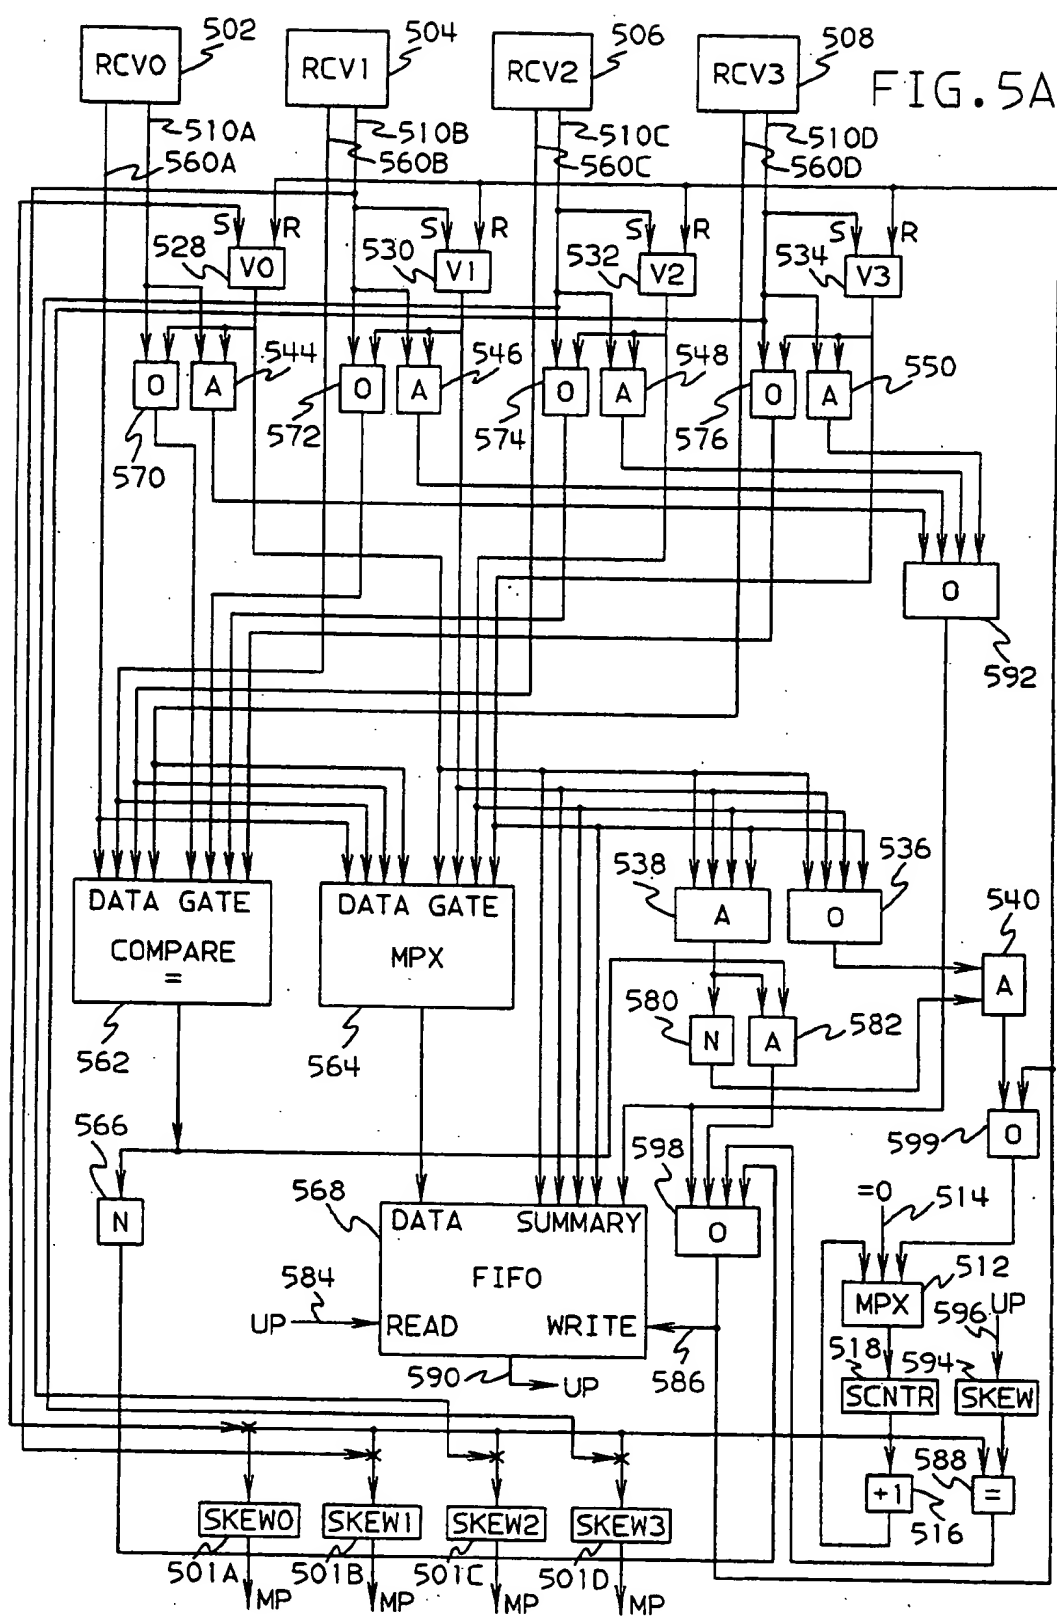



FIG. 6



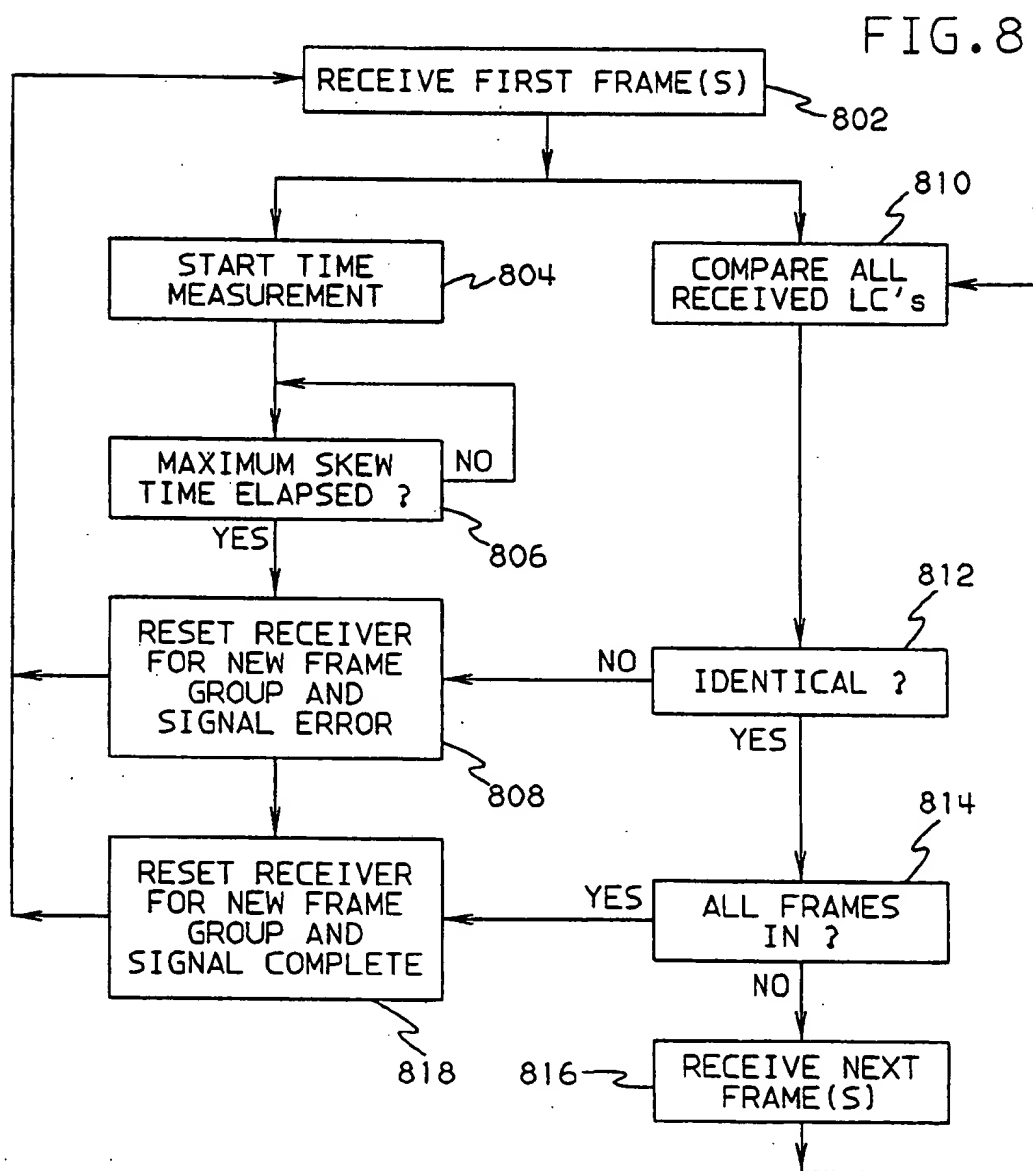
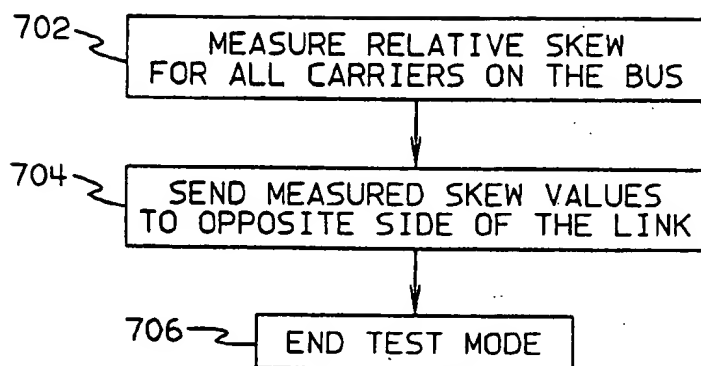




FIG. 9

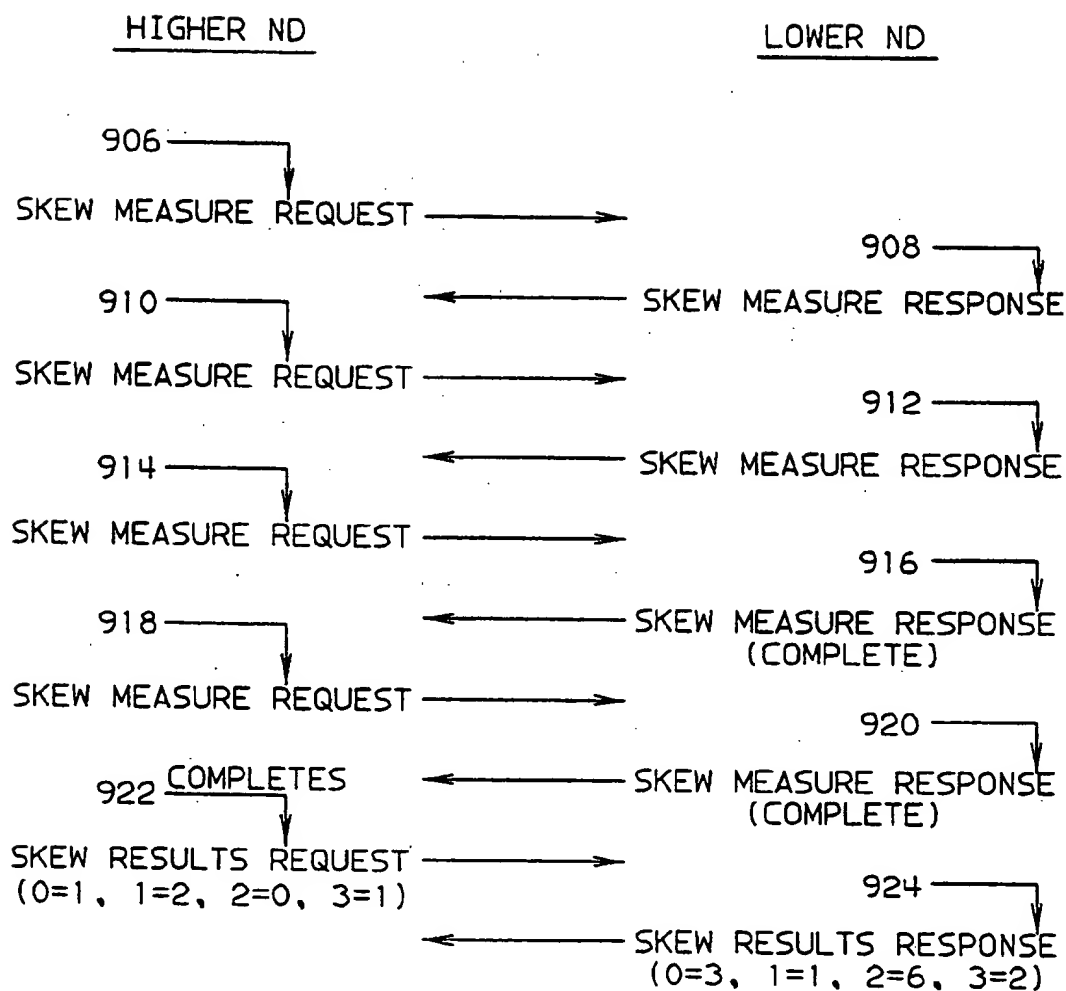


FIG. 10

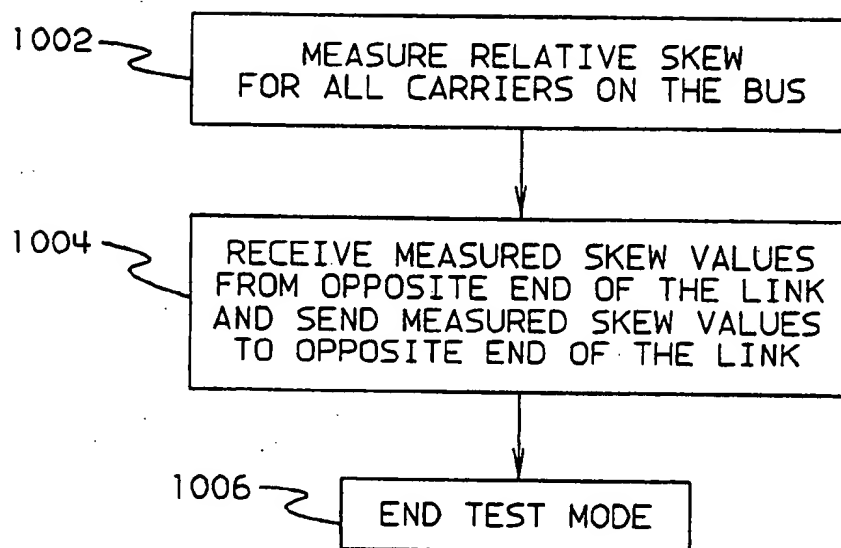


FIG. 11

1102

MIN	1104	MAX	1106	
				RECEIVER 0 (RCV0)
				RECEIVER 1 (RCV1)
				RECEIVER 2 (RCV2)
				RECEIVER 3 (RCV3)

SKEW TABLE

FIG. 12

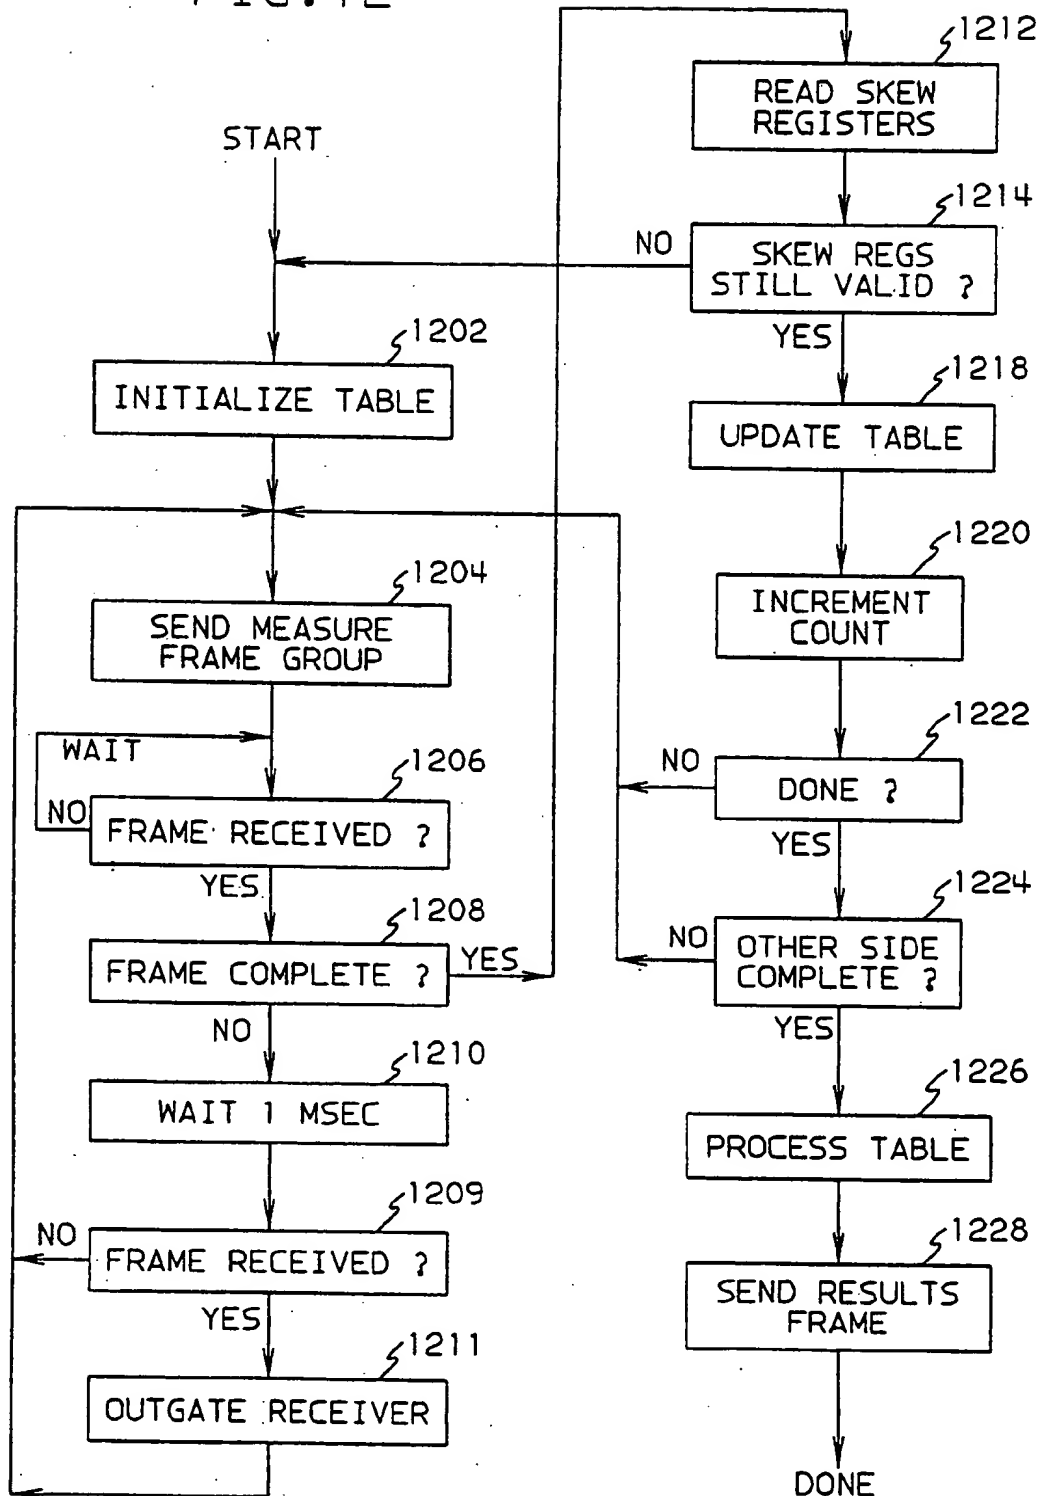


FIG. 13A


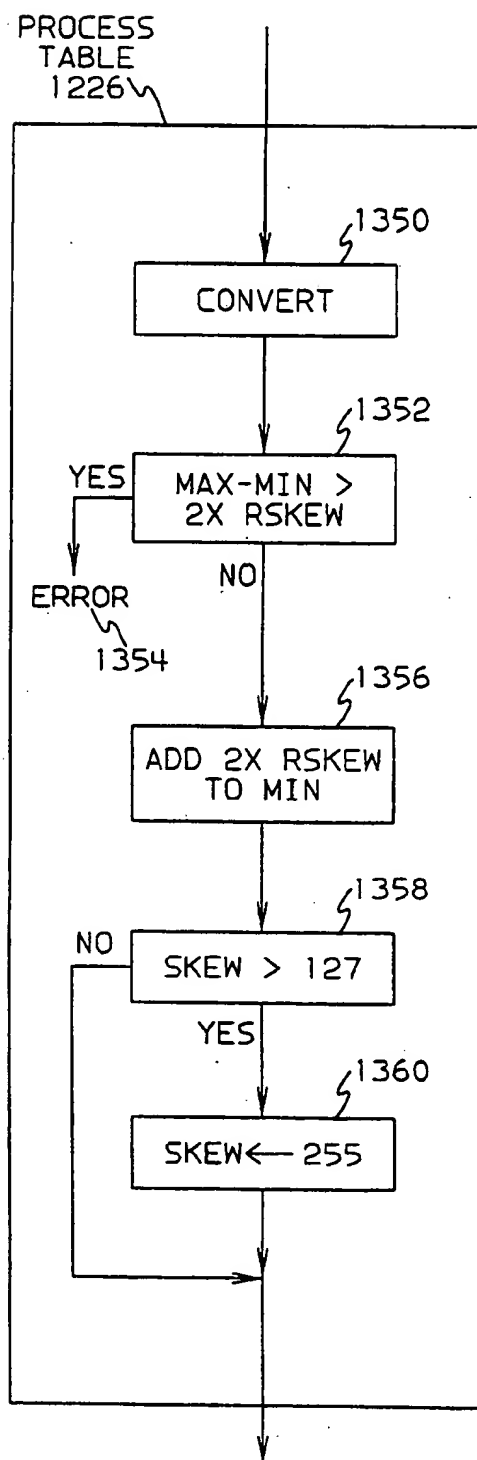


FIG. 13B



# **SKIEW MEASUREMENT FOR RECEIVING FRAME-GROUPS**

## **FIELD OF THE INVENTION**

This invention relates to computer systems, and particularly to a transceiver and method for improved skew measurement for receiving frame-groups used in highly parallel coupled systems.

## **CROSS REFERENCE TO RELATED APPLICATIONS**

This application contains subject matter which is related to the following U.S. patent applications, and while of more general applicability is an improvement which can be used in connection with these related inventions. Many of these referenced applications have been published in corresponding European and or Japanese applications as of the filing date of this application, but the numbers thereof have not been identified as of the filing of this application. They are:

- "Configurable, Recoverable Parallel Bus" by N. G. Bartow et al, Ser. No. 07/839,657, filed Feb. 20, 1992;
- "High Performance Intersystem Communications For Data Processing Systems" by N. G. Bartow et al, Ser. No. 07/839,652, filed Feb. 20, 1992;
- "Frame-Group Transmission And Reception For Parallel/Serial Buses", by N. G. Barrow, et al, Ser. No. 07/839,986, filed Feb. 20, 1992 and now U.S. Pat. No. 5,267,240 granted Nov. 30, 1993;
- "Method and Apparatus for Distributed Locking of Shared Data, Employing a Central Coupling Facility" by D. A. Elko et al, Ser. No. 07/860,808, filed Mar. 30, 1992 and now U.S. Pat. No. 5,339,427 granted Aug. 16, 1994;
- "Communicating Messages Between Processors And A Coupling Facility" by D. A. Elko et al, Ser. No. 07/860,380, filed Mar. 30, 1992;
- "Sysplex Shared Data Coherency Method and Means" by D. A. Elko et al, Ser. No. 07/860,805, filed Mar. 30, 1992;
- "Command Quiesce Function" by D. A. Elko et al, Ser. No. 07/860,330, filed Mar. 30, 1992 and now U.S. Pat. No. 5,339,405 granted Aug. 16, 1994;
- "Storage Management For a Shared Electronic Storage Cache" by D. A. Elko et al, Ser. No. 07/860,807, filed Mar. 30, 1992;
- "Management Of Data Movement From A SES Cache To DASD" by D. A. Elko et al, Ser. No. 07/860,806, filed Mar. 30, 1992;
- "Command Retry System" by D. A. Elko et al, Ser. No. 07/860,378, filed Mar. 30, 1992;
- "Integrity Of Data Objects Used To Maintain State Information For Shared Data At A Local Complex" by D. A. Elko et al, Ser. No. 07/860,800, filed Mar. 30, 1992 and now U.S. Pat. No. 5,331,673 granted Jul. 19, 1994;
- "Management Of Data Objects Used To Maintain State Information For Shared Data At A Local Complex" by J. A. Frey et al, Ser. No. 07/860,797, filed Mar. 30, 1992;
- "Recovery Of Data Objects Used To Maintain State Information For Shared Data At A Local Complex" by J. A. Frey et al, Ser. No. 07/860,647, filed Mar. 30, 1992;
- "Message Path Mechanism For Managing Connections Between Processors And A Coupling Facility"

ity" by D. A. Elko et al, Ser. No. 07/860,846, filed Mar. 30, 1992;

"Method And Apparatus For Notification Of State Transitions For Shared Lists of Data Entries" by J. A. Frey, et al., Ser. No. 07/860,809, filed Mar. 30, 1992;

"Method And Apparatus For Performing Conditional Operations on Externally Shared Data" by J. A. Frey, et al., Ser. No. 07/860,655, filed Mar. 30, 1992;

"Apparatus And Method For List Management in A Coupled DP System" by J. A. Frey, et al., Ser. No. 07/860,633, filed Mar. 30, 1992;

"Interdicting I/O And Messaging Operations In A Multi-System Complex" by D. A. Elko, et al., Ser. No. 07/860,489, filed Mar. 30, 1992;

"Method and Apparatus for Coupling Data Processing Systems" by D. A. Elko, et al., Ser. No. 07/860,803, filed Mar. 30, 1992 and now U.S. Pat. No. 5,317,739 granted May 31, 1994.

The following related applications are unpublished:

"Quiesce and Unquiesce Function for Intersystem Channels", by Neil G. Barrow, et al., Ser. No. 08/071,154, filed Jun. 1, 1993;

"Configurable, Recoverable Parallel Bus", by Neil G. Barrow et al, Ser. No. 08/071,146, filed Jun. 1, 1993;

"Frame Group Transmission and Reception for Parallel/Serial Buses", by Neil G. Barrow, et al, Ser. No. 08/071,115, filed Jun. 1, 1993;

"Concurrent Maintenance of Degraded Parallel/Serial Buses", by Kenneth J. Fredericks et al, Ser. No. 08/070,587, filed Jun. 1, 1993;

"Null Words for Pacing Serial Links to Driver and Receiver Speeds", by Daniel F. Casper et al, Ser. No. 08/071,150, filed Jun. 1, 1993.

"Error Detection and Recovery in Parallel/Serial Busses", by T. A. Gregg, et al, Ser. No. 08/070,589, filed Jun. 1, 1993.

"Dumping Service Facility for Multisystem Environments", by D. A. Elko, et al, Ser. No. 08/073,909 filed Jun. 8, 1993.

The following applications are filed concurrently herewith:

"Mechanism for Receiving Messages at a Coupling Facility", by D. A. Elko, et al, Ser. No. 08/147,697, filed Nov. 4, 1993;

"Suspending, Resuming, and Interleaving Frame-Groups", by T. A. Gregg, et al, Ser. No. 08/148,113, filed Nov. 4, 1993; and

"Frame-Group Reception and Processing", by T. A. Gregg, et al, Ser. No. 08/147,702, filed Nov. 4, 1993.

"A Message-Dispatching Mechanism for Messages Received at a Coupling Facility", by D. A. Elko, et al, Ser. No. 08/147,703, filed Nov. 4, 1993. These co-pending applications and the present application are owned by one and the same assignee, International Business Machines Corporation of Armonk, N.Y.

The descriptions set forth in these co-pending applications are hereby incorporated into the present application by this reference.

## **BACKGROUND OF THE INVENTION**

Fiber optics enable the transmission of long strings of data in a serial fashion from a driver to a receiver at long

distances (kilometers) at very high data rates (billions of bits per second). This is in contrast to traditional electrical wires which allow data to be transmitted for only short distances at these data rates (10's of meters).

High performance computers require system buses between elements of a computing system to have even higher bandwidths (100's of millions of bytes per second). The use of a single fiber optic cable running at its maximum bandwidth can be insufficient to serve as a system bus between elements of a computing system.

In the prior art, a classical solution to the problem of inadequate bandwidth on a single carrier is to transmit the bits of each data word, in parallel, one over each of a plurality of carriers (a parallel bus). Another solution, known as striping, entails sending a plurality of serial data words, in parallel, one over each of a plurality of carriers. These solutions, however, share a common problem in that skew can develop between the parallel bits or words as they are sent over the bus.

Skew is a propagation time variation between the fastest and the slowest carrier. If not handled properly, such propagation time variations can cause data errors at the receiver. Skew can be caused by differences in the physical characteristics of the driver, receiver and transmission medium for each carrier of the bus.

There are two general categories of skew. One is random skew (RSKEW) and appears as jitter. Sources of this skew are due to clock synchronization and null word insertion. The other type of skew is systemic (SSKEW) in that one transceiver is always slower or faster than another by a fixed amount of time. Differences in carrier propagation times are a typical source of the systemic type of skew.

The transmitter, carrier, and receiver each introduce skew. Below is a description of how each of these elements introduces skew.

Transmission launch skew is the time difference between the start of frame of the earliest transmitter and the start of frame of the latest transmitter. This type of skew is measured at the transmitters. There are three sources of transmission launch skew. The first is random and is caused by the synchronization of the transmitter clock to the system clock. Where the all transmitter clocks are locked to a common reference clock, and there is a common system clock synchronization point, the launch skew can be as small as one bit time, or even less. The second source of transmission launch skew is also random and is caused by null insertion. The third source of transmission launch skew is systemic and is due to differences in internal signal propagation times from transceiver to transceiver. Any reasonable physical design should limit this skew to several bit times.

A significant amount of skew is caused by the characteristic of the carrier itself. The carrier skew is simply the difference in propagation times between the slowest and fastest carriers in a link. This skew is almost entirely systemic. Slight variations can occur over time due to environmental changes such as temperature. Thus, it should be understood that even if a plurality of frames are launched at the same time, they are likely to arrive at different times because of the carrier skew.

Skew can also be caused by the characteristics of the receiver. There are two sources of skew in the receiver. The first is random and is caused by the synchronization of the system clock to the receiver clock. The second source of receiver skew is systemic and is due to differences in internal signal propagation times between different transceivers. As with the transmitter, any reason-

able physical design should limit this skew to several bit times.

To avoid data errors of the type which can be caused by skew, some means for synchronization of the parallel bits is conventionally employed. Traditional parallel bus designs use a common clock to achieve synchronization. This solution, while workable, is inefficient in that the common clock typically causes the data collection to be performed in lock step, at a speed limited at the upper end by the slowest bit path transmission time.

Thus what is needed is an improved method for transmitting data words over a plurality of carriers such as multiple fiber optic filaments.

### SUMMARY OF THE INVENTION

The present invention comprises a system and method used in asynchronously transmitting data blocks, in parallel, across multiple carriers in a serial manner. Frame groups are provided as a mechanism to transmit associated data serially on each fiber and tie the data being transmitted together. The receiver is provided with a previously measured maximum skew value for a given link from either end of the communications link. Using this information, the receiver determines which frames are part of a frame group by the arrival times of the individual frames. This aspect of the invention is achieved by a new method and a new transceiver which enables the improved method to be performed.

Thus our present invention has two aspects, a system aspect for the transceiver and a method aspect which employs the new hardware and by which the receivers on each side of the bus determines the relative skew for each conductor by performing skew measurements on a series of calibration frames generated by the transmitters on the other side of the bus. Multiple skew measurements registers and a new code algorithm provide a greatly faster method of measuring the relative skew. Each time a skew-measure frame is received, the skew information for all of the conductors is gathered at the same time. This information is then processed for all of the conductors by the code. Advantageously, this method greatly reduces the time required to perform the skew measurement. This method is fast enough that it can be used in all situations where skew measurement is required, including the more performance critical link recovery required. Additionally, the random element of the skew (RSKEW) can also be measured to determine if the system has been designed correctly.

The method can be employed in the related inventions, with a change to the transceiver receiver section, which provides a new transceiver for the links between two computing elements. This transceiver is used for asynchronously transmitting data blocks, in parallel, across multiple fibers in a serial manner. Frame groups are provided as a mechanism to transmit associated data serially on each fiber and tie the data being transmitted together. The frame groups do not have sequence numbers, therefore, the receiver determines which frames are part of a frame group by the arrival times of the individual frames. The transceivers for each member of the parallel bus asynchronously achieve synchronism from either end of the fiber. A receiver of the transceiver determines which frames are part of a group and determines a relative skew on a conductor of a link by performing skew measurements on a series of calibration frames sent by a transmitter on another end of said link. A skew monitor performs skew measurements and includes a plurality of skew measurement registers for

measuring the relative skew between frames of a frame group, each register having an input from a skew counter, and outputs coupled to a microprocessor controlling the monitor. Each register gates in the value in the skew counter at the time a link control word is received. The receiver can determine that one or more frames of the calibration frames have been lost and have not exceeded the maximum measurable skew limits.

According to an embodiment of the present invention, the transceivers for each member of the parallel bus asynchronously achieve synchronism at each end of the fiber. Thus the need for a common clock is eliminated. The receivers on each side of the bus determine the relative skew for each conductor by performing skew measurements on a series of calibration frames generated by the transmitters on the other side of the bus. When the skew on all conductors, viewed from both sides of the bus, has been determined, the skew values are exchanged across the bus, thus enabling the transmitters to set proper frame spacing.

Advantageously, skew measurement according to the present invention enables data to be sent across a link comprising a multi-carrier bus, limited only by the maximum bandwidth of the link itself rather than by the worst case theoretical skew. According to an embodiment of the present invention, data integrity is ensured by first measuring the relative skew values for the carriers in a link, at the time that the link is established or reinitialized. After the relative skew values have been determined, they are exchanged between both sides of the link and a connection is established having the start-to-start spacing of transmitted frames limited to the worst case (largest) skew value measured by the receivers. Since frames can be made very large, a significant amount of data can be sent across the link, at the maximum link bandwidth without concern for skew. Further, where large data frames are used, the bandwidth limiting effect of limiting the start-to-start frame spacing to the worst case skew value for a given link can be kept at a minimum.

It will be appreciated that the present inventions of a transceiver and method for improved skew measurement for receiving frame-groups used in highly parallel coupled systems represents an improvement to the system disclosed in the continuation-in-part of U.S. patent application Ser. No. 07/839,986 filed Feb. 20, 1992 entitled "Frame Group Transmission and Reception for Parallel/Serial Buses", the continuation application being U.S. Ser. No. 08/071,115, filed Jun. 1, 1993.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be better understood by reference to the drawings in which:

FIG. 1 is a block diagram of a physical link between two computing elements;

FIG. 2 is an illustration of the format of an exemplary frame;

FIG. 3 is an illustration of the format of an exemplary frame group;

FIG. 4 is a logic diagram of an exemplary transmitter;

FIG. 5 is a logic diagram of an exemplary receiver;

FIG. 5A is a logic diagram of an exemplary receiver with additional skew measurement registers, illustrating our improved transceiver receiver.

FIG. 6 is a flow diagram of a skew test procedure;

FIG. 7 is a flow chart summary of the skew test/measurement procedure as view from a receiving end of the link;

FIG. 8 is a flow chart of frame reception and frame group integrity checking in a receiver according to an embodiment of the present invention;

FIG. 9 is a flow diagram of an alternative skew test procedure. Like reference numerals appearing in more than one FIGURE designate like elements;

FIG. 10 is a flow chart summary of an alternative skew test/measurement procedure as view from a receiving end of the link;

The transceiver receiver in monitoring and processing skew in the process of skew measurement utilizes a new method employed in the transceiver receiver, as illustrated by FIGS. 11-13B. In these FIGURES:

FIG. 11 is a diagram of the skew table used in the improved skew measurement procedure;

FIG. 12 is a flow chart of the improved skew measurement procedure;

FIGS. 13A and 13B show detailed flowcharts of the skew table update and the skew table final processing used in the improved skew measurement procedure.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Before describing particularly the improvements which result in a new transceiver and method for skew measurement, the environment of the related applications will be described.

Turning first to FIG. 1, a physical link between two computing elements 102, 104 is illustrated. These elements could be, for example, two computers or a computer and a shared memory device. In any event, the computing elements 102, 104 are connected by way of an intersystem channel link 106 comprising a fiber optic bus 108. The fiber optic bus 108 is formed of multiple fiber pairs 110. Each fiber pair consists of two optical fibers, one for transmitting information and one for receiving information. While any number of fiber pairs can be used, exemplary values for intersystem channels would be a range of 1 to 64 fiber pairs in powers of two.

The fiber pairs 110 of the fiber optic bus 108 are coupled to the computing elements 102, 104 by way of transceivers 112, 114 located at opposite ends bus. Each of the transceivers 112, 114 includes a transmitter unit and a receiver unit, both of which will be described in more detail later.

Data is sent across the link from an originator buffer 116 in a transmitting computing element. The received data is then collected and temporarily stored in a recipient buffer 118 in a receiving computing element. Although one buffer is shown in each of the computing elements of FIG. 1, it should be understood that a computing element can dynamically set up various numbers of originator and/or recipient buffers depending on the number and type of links to be established. The buffers 116, 118 can be of the "store-thru" type, meaning that they need not be filled before they are unloaded. This type of buffering has the advantage of lower transmission and reception latency times.

The information to be sent is divided up among the fiber pairs 110 that form the fiber optic bus 108. The information is sent from the originator buffer 116 to the recipient buffer 118 in the form of frames. Each frame is transferred on a single fiber of the bus 108. More information on configuration of the link and the structure of the buffers are contained, respectively, in copending applications entitled CONFIGURABLE, RECOVERABLE PARALLEL BUS to Barrow et al. (Ser. No. 07/839,657) and HIGH PERFORMANCE CHAN-

NELS FOR DATA PROCESSING SYSTEMS to Barrow et al. Ser. No. 07/839,652), both filed on Feb. 20, 1992, referenced above.

The format of an exemplary frame is illustrated in FIG. 2. Each frame starts with a link-control word (LC) 202. Various fields in the link-control word 202 identify the frame format and type, designate a buffer area, and control the state of the transceiver and link.

A link-control-CRC (cyclical redundancy check) word 204 follows the link-control word. The link-control CRC word 204 is conventionally generated from the values in the link-control word. The link-control CRC word is checked at the receiver to test the validity of the link-control word in the incoming frame.

There are two types of frames, control frames and information frames. Control frames do not require an information field. They consist only of a link-control word and a link-control CRC word. An information frame has a link-control word 202 a link-control CRC word 204 and an information field 206. Information fields contain, for example, from one to 1,024 words. The information field is used to send information from a buffer area at one end of the link to a buffer area at the other end.

An information field is followed by an information-field CRC word 208. The information-field CRC word is conventionally generated from the values in the information field. The information field CRC word is checked at the receiver to test the validity of the information field in the incoming frame.

Related information can be placed in several frames, one on each transceiver of the operational link. These frames (referred to as a "frame group") are transmitted concurrently. There can be as many frames in the group as there are transceivers in the operational link. FIG. 3 shows an exemplary frame group 302. The use of the frame group enables the data to be sent at a bandwidth that is multiples of the single fiber optic bandwidth since multiple frames (one per fiber) are transmitted simultaneously across the link.

Although not provided with sequence numbers, the link control words in each frame group are interlocked with those of subsequent and previous frame groups. The interlocking can be accomplished in a number of different ways. For example, each transmission may start with a request, followed by data, followed by an acknowledge. Each of these types of transmissions having a unique link control word. Further, a sending and/or recipient buffer address can be sent within the link control words. Thus data transmissions from (and/or directed to) different buffers will have different link control words.

In the preferred embodiment, frames associated with the same buffer and the same type of frame group may have the same link control words. For example a given message may include 16 data frames, 15 of which will have the same LC and one of which will have a distinct LC in that it includes a bit indicating the start of a buffer area. In this case, a lost frame among the 15 having an identical LC can not necessarily be detected by comparing the LCs in each received frame group. Instead, either the comparison of LCs in the last frame group will detect the problem because a frame belonging to the next frame group (of a different type and having a different LC) has been received, or the skew timer will expire (as described later). In either event, the receiver will detect that all data has not been received for the buffer area. This is accomplished by comparing the total

number bytes actually received with the total transmitted byte count which is included as part of every LC.

As an alternative embodiment, the transmitter can ensure that the link control words in each frame group are identical with each other, but are different than those in the immediately previous and subsequent frame groups.

Information from an originator buffer is sent across the link in one or more frame groups. When several frame groups are used to transmit the contents of one buffer, the frame groups may be sent in a sequence corresponding to the ascending order of the words in the buffer. For example, groups of lower numbered words are sent before groups of higher numbered words. When the operational link comprises more than one transceiver, the buffer contents are sent in units of words on a transceiver basis. That is to say, the words are distributed across the transceivers (and thus the fibers), modulo the number of transceivers in the link (e.g. transceiver 0 has word 0, 4, 8, 12 in a four transceiver operational link).

According to an embodiment of the present invention, the transceivers for each carrier of the bus asynchronously achieve synchronism at each end of the fiber. Thus the need for a common clock is eliminated. Various means for synchronizing frames transmitted as serial data over a signal conductor are known in the art and thus will not be discussed in detail here. An example of such synchronization means is described in U.S. Pat. No. 4,970,714 to Chen et al. and assigned to the same assignee as the present invention.

Each fiber and its associated transmission and reception hardware has differing propagation characteristics which result in skew. Skew is the propagation time variation between the fastest and the slowest fiber. Reception of a start of a frame on any fiber causes the receiver to start checking for reception of the complete frame group and to reassemble the received frames of data in the recipient buffer. The receiver determines which frames are part of a frame group by the arrival times of the individual frames and the contents of the link control words. If all of the frames of a frame group are not received within an established skew window or if the link control words received within the skew window are not identical, a skew error is recognized.

The skew rules, according to an embodiment of the present invention, set the skew window in which all frames of a frame-group must be received in order to receive a valid message. The rules also set the rate at which consecutive frame groups can be transmitted so that frames of one frame-group are not confused with frames of another frame-group. Further, examination of link control words for consistency within a given frame group protects the system from "lost frame" effects.

For a given link, the skew of frame groups measured at the receiver varies by plus or minus the random skew (RSKEW) component. The maximum RSKEW component is set by the design of the link hardware while the systemic skew (SSKEW) component is largely due to the optical cable installation. According to an embodiment of the present invention, a measurement of the total skew (the optical and random components) is made at the time the link is initialized to ensure that it does not exceed set limits and further to ensure that frames belonging to a common frame group are properly identified as such by the receiver.

An exemplary transmitter unit according to an embodiment of the present invention is illustrated in FIG.



4. A four transceiver link is shown. The individual transmitters (one for each fiber optic carrier) are shown as XMIT0 402, XMIT1 404, XMIT2 406, and XMIT3 408. It should be understood that the transmitter of FIG. 4 can be operated using fewer than four carriers (e.g. 3, 2 or 1) by gating out selected transmit portions.

The internals of XMIT0 402 are shown. Along with the data paths (not shown), each transmitter has an input named START 410 which starts frame generation. Each transmitter also has an output called READY 412 which indicates that the current frame transmission has ended, and that the skew wait time has elapsed.

Within each transmitter there is a WAIT register 414 which is initialized at link initialization time. The value in the WAIT register 414 is the number of words which must be transmitted between LC (Link Control) words. The Wait register 414 is connected to a 3 input multiplexor (MPX) 416. The output of the MPX 416 is connected to a wait counter register (WCNTR) 418. The output of the WCNTR 418 is connected to decremter (-1) 420, an 8 input nand circuit (=0) 422 detecting a value of zero in the WCNTR 418, and back to a first one (A) of the MPX 416 inputs (A-C). The output of decremter 420 is connected to the third data input (C) of the MPX 416.

The channel starts a frame when all four transmitters are ready (i.e. all transmitters are idle and the minimum LC word to LC word spacing time has elapsed) by turning on the START line 410 for 1 cycle. This line causes each transmitter to begin frame transmission. When the Outbound Frame State Machine (OFSM) 424 sends the LC (indicated on line 426), the END 428 and MET 430 Shift Register Latches (SRLs) are reset, and the value in the WAIT register 414 is gated through the MPX 416 and the WCNTR 418. As words are sent on the link, the Data Request line (DR) 432 is active, and the WCNTR 418 is decremented via the decremter 420 and MPX 416 circuits. The End Of Frame (EOF) may become active before, during, or after the WCNTR 418 reaches the value of zero. The EOF line 434 sets the END SRL 428, and the 8 input nand circuit 422 sets the MET SRL 430. When both SRLs 428, 430 are on, the output of the AND circuit 436 is active, and all conditions have been satisfied for the start of the next frame for that transceiver.

The READY lines 412 from the transmitters set the RDY0 438, RDY1 440, RDY2 442, and RDY3 444 SRLs. The outputs of these SRLs are connected to a 4 input AND circuit 446. When all 4 RDY SRLs are set, the output of the AND circuit 446 indicates that all 4 transmitters are ready to start the next frame. The output of this AND circuit 446 is also used to reset the RDY SRLs 438-444.

An exemplary receiver unit according to an embodiment of the present invention is illustrated in FIG. 5. A four transceiver link is shown. The receivers are shown as RCV0 502, RCV1 504, RCV2 506, and RCV3 508. The data paths 560 (A-D) from each receiver contain the current LC detected by the receiver. The output lines 510 (A-D) are active for 1 cycle each time a new LC is detected by the corresponding receiver. The contents of the data paths 560 (A-D) are compared by the gated compare logic (COMPARE) 562, and are also selected by the gated multiplexor (MPX) 564.

The Set-Reset latches (S-R latches) V0 528, V1 530, V2 532, and V3 534 are set by the respective lines 510 (A-D) from the receivers. When a receiver detects a

new LC, the corresponding S-R latch 528-534 is set. The four S-R latches 528-534 control the MPX 564, and are inputs to the first-in first-out buffer (FIFO) 568, a four input AND gate (A) 538, and a four input OR gate (O) 536. The four OR gates 570-576 each have two inputs. One input is a line 510 (A-D) from a corresponding one of the receivers, and the other input is the output of a corresponding one of the S-R latches 528-534. The output of these OR gates 570-576 signal to the COMPARE 562 and determines which of the data paths to compare.

The skew measurement function is performed by a counter (SCNTR) 518 which is fed from a second multiplexor (MPX) 512. The second multiplexor 512 gates either a zero value 514 or the incremented value (+1) 516 to the skew counter (SCNTR) 518. The SCNTR 518 is an eight bit register. The output of the SCNTR 518 feeds an equal-to compare circuit (=) 588 and the incrementer (+1) 516. The other input to compare 588 is from the output of a register (SKEW) 594. When the value in SCNTR 518 is equal to the value in SKEW 594, the output of the compare circuit 588 is active. The value in SKEW 594 is set by a microprocessor (not shown) over line 596.

The FIFO 568 stores an LC selected from one of the receivers along with summary information describing which of the receivers detected the LC. The write line 586 to the FIFO 568 is controlled by a four input or (O) circuit 598. The microprocessor reads entries in the FIFO 568 over a data bus 590 by activating the read line 584.

As with the transmitter unit, receivers for any of the carriers can be gated out of the system under control of the microprocessor via a control register (not shown).

Operation of an exemplary receiver unit according to an embodiment of the present invention as illustrated in FIG. 5 starts with the receipt of an LC by at least one of the receivers and finishes with a write operation to the FIFO 568. When the first LC is received by one or more of the receivers RCV0 502, RCV1 504, RCV2 506, and RCV3 508, the corresponding S-R latch 528-534 is set. The output of OR circuit 536 is activated which starts the increment of the SCNTR 518. When all four LC's are received, all four S-R latches 528-534 are set and the output of AND gate 538 is active (HIGH). This output feeds an inverter 580 which, in turn, feeds an AND gate 540. Thus, when AND gate 538 is active (HIGH) the output of AND gate 540 is deactivated (LOW) which, in turn, stops the SCNTR 518 through AND gate 540 and OR gate 599. If all LC's have been received and they all compare equal (as detected by COMPARE 562) the output of AND gate 582 is active. This output feeds OR gate 598 which, in turn, causes a write into the FIFO 568. Along with the LC from the MPX 564, summary information describing which receivers have an LC is also stored in the FIFO 568.

If an LC is damaged by transmission noise on the link, it is not detected by the receiver. When this happens the output of AND circuit 538 does not activate, and the SCNTR continues to count until its value reaches the value in SKEW 594. At this point, the output of the compare circuit 588 becomes active and it feeds OR circuit 598 which causes an entry to be made into the FIFO 568. In this situation, the summary information in the FIFO 568 indicates which of the receivers did not receive an LC.

Another possible result of a damaged LC is the apparent overlap of two different frame groups. Suppose that

two frame groups reach the receivers with the minimum LC to LC spacing, and suppose that the first frame of the first group is damaged and discarded. In this case, the activation of the output of OR gate 536 is delayed until the next LC of the first frame group is received. This delay also starts the SCNTR 518 later. After all of the undamaged LC's have been received, the SCNTR 518 is still running. At this point the first frame of the second frame group may arrive. If the LC of this frame is different than the other LC's at the input of COMPARE 562, the output of COMPARE 562 will activate (since there is a non-compare of the LC's indicating that the LC's are from two different frame groups). An inverter (N) 566 detects the non-compare and feeds OR gate 598 which, in turn, causes an entry into the FIFO 568. The summary information indicates the LC and the receivers which detected this LC. The action of OR gates 570-576 bypasses the delay caused by the S-R latches 528-534 so that the LC which caused the non-compare is not included in the summary information.

OR gate 592 receives inputs from AND gates 544-550. Each of the AND gates 544-550 detects if its respective receiver receives an LC before the previous LC (from the same receiver) has been entered into the FIFO 568. Specifically, the AND gates 544-550 detect when a new LC is received (as indicated by lines 510A-D respectively) before the S-R latches 528-534 have been reset. This situation can occur only if the transmitter feeding the receivers does not obey the skew rules established at initialization time. The output of OR gate 592 feeds OR gate 598 which causes an entry to be made into the FIFO 568. The summary information includes the output of OR gate 592.

In all of the above situations, each time a FIFO 568 entry is made, the S-R latches 528-534 are reset. In the case of the LC non-compare, the S-R latch for the receiver which caused the non-compare is set even though the reset line is active. This is so because the S-R latches' set function overrides the reset function. Also, each time a FIFO 568 entry is made, the SCNTR 518 is reset to zero by the action of OR gate 599 and MPX 512.

It is noted that the transmitters 402-408 of FIG. 4 are not synchronized by a common clock. Similarly, the receivers 502-508 are also not synchronized with one another by a common clock. Instead, each time the link is established the transmitter and receiver pairs achieve synchronization with each other in an asynchronous manner, just as if they connected by the only carrier in the link. Further, the receivers 502-508 use reception of the frames themselves, the skew rules and LC checking as described above to ensure that frames within the same frame group are properly identified as such. Advantageously, this eliminates the need for synchronizing hardware.

A flow diagram of a skew test/measurement procedure according to an embodiment of the present invention is illustrated in FIG. 6.

When skew measurements are to be performed, any transceiver may send a skew-test-request (STR) control frame 602 (or a replicated STR control frame) out on the link. By replicated, it is meant that an identical copy of a given frame (e.g. the STR frame) is transmitted on each carrier within a 2.5 word time transmit window. The skew-test-request frame includes a link control word identifying it as an STR. Nodes on both sides of the link may send out a skew test request at the same time. In this case, the node with the higher address

becomes the master (i.e. controls the skew measurement procedure) and the node with the lower address becomes the slave. If only one node has sent the STR frame, it becomes the master and the recipient becomes the slave notwithstanding the relative node address.

Once a skew-test-request frame is received by any one of the receivers of the slave receiver unit (and the master slave relationship is established), the slave receiver unit places itself into a skew measurement mode (thereby causing its internal logic to respond as described below), resets its skew counter value to zero, and sets the skew register 594 to a predetermined initial value (e.g. 255). It should be understood, however, that another initial value could alternatively be chosen based on actual experience with a given system. The slave node then sends back a replicated skew-test-response control frame 604. This frame indicates to the master that the slave is ready to begin the skew measurement procedure.

In response to receipt of a skew-test-response frame, the master receiver unit also places itself into skew measurement mode, resets its own skew counter to zero, loads its skew register 594 with an initial value (e.g. 255) and sends out a replicated skew-measure-request (SMR) control frame 606.

When the first slave receiver(s) 502-508 receive the first skew-measure-request frame, the slave receiver unit determines whether all of the frames have arrived within the skew value indicated by the skew register 594. If not, the receiver unit's microprocessor checks the summary information in the FIFO 568 to determine which frame(s) were late (i.e. not received within the skew value indicated by the skew register 594). The receivers detecting the late frames are gated out and are not used in subsequent skew tests. The identity of these receivers and the skew register value is stored in the microprocessors local storage for later use. Next, a new value (e.g. 254) is loaded into the skew register and then the slave sends a replicated skew-measure-response control frame across the link on all carriers (including those whose receivers have been gated out).

The skew-measure-response frame has the same effect on the master receiver unit as the skew-measure-request frame has on the slave receiver unit. Specifically, the master determines which frames were received within the skew value indicated by its skew register 594, outgates the late receivers and loads a new skew value in its skew register. When this has been accomplished, the master sends a replicated skew-measure-request frame across the link on all carriers.

In response to this skew-measure-request frame, the slave receiver unit determines whether all of the frames (received on receivers that have not been gated out) have arrived within the new skew value indicated by the skew register 594. If not, the slave receiver unit's microprocessor again checks the summary information in the FIFO 568 to determine which frame group(s) were late. The receivers detecting the late frames are gated out and are not used in subsequent skew tests. Also, the identity of the gated out receivers is also stored in the microprocessors local storage for later use. Next, a new value (e.g. 253) is loaded into the skew register and then the slave sends a replicated skew measure response control frame across the link on all carriers (including those whose receivers have been gated out).

The above described process continues iteratively (e.g. 614) until a relative skew value has been deter-

mined for each receiver on both sides of the link (i.e. until all but one receiver has been gated out or the skew register has been decremented to zero). When the slave has measured the relative skew for all receivers, it sets a bit in the skew measure response frame 616 indicating that the measurement is complete. When the master determines that skew measurements on both sides of the link are complete, it sends a skew-results-request (SRR) control frame 622 across the link. In response, the slave sends a skew results response control frame 624 (one on each carrier) back across the link. These frames include the measured skew values for the slave receivers (e.g. carrier 0=1, carrier 1=0, carrier 2=6, carrier 3=2). Each of the skew values is indicative of a number of word transmit times (e.g. 2=2 word transmit times). It should be understood that the skew measurements are all made relative to the fastest carrier(s). Thus the fastest carrier(s) (i.e. the first to arrive across the link) will always have a relative skew value of zero. The skew value returned in the skew results response may also include a padding value added to each measurement. This padding value can be used to compensate for random skew (jitter).

Once the skew-results-response frame has been received, the master can send an establish-transceiver-group (ETG) control frame 626. This request is used by the originating (master) node to elicit responses from the set of transceivers at the recipient (slave) node and also to set the intended link (the set of carriers that will be used to establish communications) at the slave. The request is sent as a replicated frame on all of the carriers that have passed the skew test (e.g. have a relative skew value of 254 or less). The parameter field of the ETG request contains the largest of the skew values determined by the master receivers. The master can also add a pad value in a similar manner as the slave (in sending the skew results response). This value determines the link-control-word to link-control-word spacing that is required for any frame groups that are to be sent by the slave.

In response to the ETG request, the slave sends an ETG response frame 628 (which is replicated on carriers on which the ETG request was received). The master node checks the ETG response to ensure that it has been received on each of the carriers on which the ETG request was sent. If the master node receives the ETG responses on all transceivers of the intended link, the link is established using the carriers designated by the ETG frames. If not, a link can be retried or established using a subset of the carriers comprising those which are operational. In that event, the master sends another skew results request to the slave and the slave again responds with the previously measured skew. Only the maximum skew on the intended link carriers is used to determine the minimum frame spacing.

A flow diagram of an alternative skew test/measurement procedure according to an embodiment of the present invention is illustrated in FIG. 9. Note that ND is Node Descriptor. This procedure enables measurement of skew from either end of the communications link.

When skew measurements are to be performed, a transceiver may send a skew-measure-request (SMR) control frame 906 (or a replicated SMR control frame) out on the link. By replicated, it is meant that an identical copy of a given frame (e.g. the SMR frame) is transmitted on each carrier within a 2.5 word time transmit window. The skew-measure-request frame includes a

link control word identifying it as an SMR. Only the node that is the master may send out a skew measure request. The master may be determined by being the node with the higher address (i.e. controls the skew measurement procedure) and the node with the lower address becomes the slave, or may be assigned by design conventions.

The first time that a skew-measure-request frame is received by any one of the receivers of the slave receiver unit (and the master slave relationship is established), the slave receiver unit places itself into a skew measurement mode (thereby causing its internal logic to respond as described below), resets its skew counter value to zero, and sets the skew register 594 to a predetermined initial value (e.g. 255). It should be understood, however, that another initial value could alternatively be chosen based on actual experience with a given system. The slave node then sends back a replicated skew-measure-response control frame 908. This frame indicates to the master that the slave is ready to begin the skew measurement procedure.

When the master unit sends the first skew-measure-request frame, the master receiver unit also places itself into skew measurement mode, resets its own skew counter to zero, loads its skew register 594 with an initial value (e.g. 255) and sends out the first replicated skew-measure-request (SMR) control frame 910.

When the first slave receiver(s) 502-508 receive the next skew-measure-request frame, the slave receiver unit determines whether all of the frames have arrived within the skew value indicated by the skew register 594. If not, the receiver unit's microprocessor checks the summary information in the FIFO 568 to determine which frame(s) were late (i.e. not received within the skew value indicated by the skew register 594). The receivers detecting the late frames are gated out and are not used in subsequent skew tests. The identity of these receivers and the skew register value is stored in the microprocessors local storage for later use. Next, a new value (e.g. 254) is loaded into the skew register and then the slave sends a replicated skew-measure-response control frame across the link on all carriers (including those whose receivers have been gated out).

The skew-measure-response frame has the same effect on the master receiver unit as the skew-measure-request frame has on the slave receiver unit. Specifically, the master determines which frames were received within the skew value indicated by its skew register 594, outgates the late receivers and loads a new skew value in its skew register. When this has been accomplished, the master sends a replicated skew-measure-request frame across the link on all carriers.

In response to this skew-measure-request frame, the slave receiver unit determines whether all of the frames (received on receivers that have not been gated out) have arrived within the new skew value indicated by the skew register 594. If not, the slave receiver unit's microprocessor again checks the summary information in the FIFO 568 to determine which frame group(s) were late. The receivers detecting the late frames are gated out and are not used in subsequent skew tests. Also, the identity of the gated out receivers is also stored in the microprocessors local storage for later use. Next, a new value (e.g. 253) is loaded into the skew register and then the slave sends a replicated skew measure response control frame across the link on all carriers (including those whose receivers have been gated out).

The above described process continues iteratively (e.g. 914) until a relative skew value has been determined for each receiver on both sides of the link (i.e. until all but one receiver has been gated out or the skew register has been decremented to zero). When the slave has measured the relative skew for all receivers, it sets a bit in the skew measure response frame 916 indicating that the measurement is complete. When the master determines that skew measurements on both sides of the link are complete, it sends a skew-results-request (SRR) control frame 922 across the link. In response, the slave sends a skew results response control frame 924 (one on each carrier) back across the link. These frames include the measured skew values for the slave receivers (e.g. carrier 0=1, carrier 1=0, carrier 2=6, carrier 3=2) and master receivers (e.g. carrier 0=1, carrier 1=2, carrier 2=0, carrier 3=1). Each of the skew values is indicative of a number of word transmit times (e.g. 2=2 word transmit times). It should be understood that the skew measurements are all made relative to the fastest carrier(s). Thus the fastest carrier(s) (i.e. the first to arrive across the link) will always have a relative skew value of zero. The skew value returned in the skew results may also include a padding value added to each measurement. This padding value can be used to compensate for random skew (jitter). This value determines the link-control-word to link-control-word spacing that is required for any frame groups that are to be sent by the master.

Once the skew-results-request frame has been received, the slave tests the skew results to determine the set of transceivers that are within the maximum allowed skew at both ends of the link and comprise an allowed intended operational link. The set of transceivers that meets this test becomes the operational link. The slave can indicate non-acceptance of an operational link by a bit in the parameter field.

Once the skew-results-response frame has been received, the master tests the skew results to determine the set of transceivers that are within the maximum allowed skew at both ends of the link and comprise an allowed intended operational link. The set of transceivers that meets this test becomes the operational link. The slave can indicate non-acceptance of the operational link by a bit in the parameter field.

The parameter field of the SRR request contains the largest of the skew values determined by the master receivers. The master can also add a pad value in a similar manner as the slave (in sending the skew results response). This value determines the link-control-word to link-control-word spacing that is required for any frame groups that are to be sent by the slave.

Advantageously, the skew measurement method according to an embodiment of the present invention can determine the total skew including the systemic (SSKEW) and random (RSKEW) elements thus reducing or eliminating the need for pad values depending on the confidence level required. In such an embodiment, the operation of each carrier at a given skew value (in the skew register 594) is verified by repeating the skew test a multiplicity of times (e.g. 100) at that value. If a carrier can not pass the skew test even once (i.e. a replicated frame is not received within the tested skew window), it is considered to not be within the tested skew window as indicated by the skew register.

A summary of the skew test/measurement procedure as viewed from a receiving end of the link is illustrated in FIG. 7 and FIG. 10. Once the receiver unit is in skew

measurement mode, in steps 702 and 1002 it measures the relative skew for all carriers on the bus. Once the measurements have been accomplished on both sides of the link, in steps 704 and 1004 the measured worst case skew values are exchanged across the link. In the case of a slave node, the measured skew values for each carrier are sent across the link to the master. Finally, in steps 706 and 1006, once the exchange is complete, the receiver unit ends skew test/measurement mode and awaits a command from the system or link.

A summary of the method used by the receiver unit to receive and verify a frame group is illustrated in FIG. 8. In step 802 the receiver unit receives the first frame or frames in a frame group. In response, in step 804 the receiver unit starts its skew counter to measure elapsed word times and continuously compares the elapsed time against the maximum measured skew value for the established link. If at any time before all frames have been received (or an error has been detected), the elapsed time matches the worst case skew value (as detected in step 806), in step 808 the receiver unit is reset to receive a new frame group (which includes resetting the skew counter) and an error is reported.

In parallel, in step 810, as frames arrive their link control words are compared to ensure that they are all identical. If any newly arrived link control word does not match those for the other received frames (detected in step 812) in step 808 an error is signaled and the receiver unit is reset to receive a new frame group. The non-comparing frames are considered to be part of the new frame group and are treated by the receiver as such.

If the newly arrived frames compare, a test is made in step 814 to determine whether all the frames in the frame group have been received. This is a matter of comparing the active carriers, as indicated by the receiver unit's control register, with the received and properly comparing frames. If all frames have not yet been received, in step 816 the receiver continues to wait for additional frames to be received and continues the comparison and time measurement process as described above. If all frames have been received before the worst case measured skew time has elapsed, the receiving end of the link signals complete and the receiver is reset to receive the next frame group.

In summary, according to an embodiment of the present invention, the skew rules are as follows:

All transceivers must start transmission of the link-control word (LC) within 2.5 word transmission times across all of the drivers sending a frame group. A word transmission time is the amount of time it takes to launch (transmit on to the link) a word of information (32 data bits plus 8 bits of encoding information in the present embodiment).

Start-to-start spacing of link-control words of back-to-back frames on a single driver may be no closer than the value that is established during the skew test procedure. The start-to-start frame spacing is contained in the establish-transceiver-group request or the largest skew value returned in the skew results response frames. Alternatively, the start-to-start frame spacing is the largest skew value for the operational link contained in the skew results request or the largest skew value for the operational link returned in the skew results response frames.

If a link control word is received on a single receiver with less separation from the previous link control word than the start-to-start frame spacing as determined

above, the frame is discarded as described by the function of OR gate 592.

Transceivers check to see that link-control words for each frame group are received within a predetermined time span from arrival of the first frame on the link. The length of the time span is specified by the maximum measured (worst case) skew value as determined by the skew test and can be optionally padded with a model dependant value that allows for long term jitter. The checking is performed only on the transceivers that are the members of the intended link as defined by the ETG request frame group sent or received and recognized. If the alternative method is used, the checking is performed only on the transceivers that are the members of the intended link as defined by the skew results request and skew results response.

If a frame group is received with a skew error, the frame group is rejected, or discarded as appropriate. Transceivers that fail a skew check are logged out for maintenance purposes.

In accordance with the improvements described herein, FIG. 5A illustrates the additional registers used by the transceiver receiver in monitoring skew. The process of skew measurement utilizes a new method employed in the transceiver receiver, as illustrated by FIGS. 11-13.

The improved method of measuring skew required development of new hardware, starting with adding SKEW registers (SKEW0 501A, SKEW1 501B, SKEW2 501C, and SKEW 501D) to the hardware shown in FIG. 5. These SKEW registers are used to measure the relative skew between frames of a frame group and are shown in FIG. 5A. Each of these registers has a single input from the SCNTR 518 whose operation is described earlier. The outputs of the SKEW 501 (A,B,C,D) registers are connected to the microprocessor. The gating in of the SKEW 501(A,B,C,D) registers is controlled by lines 510(A,B,C,D), respectively. Each time a new LC is received by RCV0 502, RCV1 504, RCV2 506, and RCV3 508, lines 510(A,B,C,D), respectively, are active for 1 cycle. As a result, the corresponding SKEW 501 (A,B,C,D) register gates in the value in the SCNTR 518 register at the time the LC is received.

The operation of the SKEW 501(A,B,C,D) registers is as follows. Before the first frame of a frame group is received, the SCNTR 518 has already been reset to zero. When the first frame of the frame group is received, the corresponding SKEW 501 (A,B,C,D) register receives the zero value from the SCNTR 518. As subsequent frames are received, their corresponding SKEW 501 (A,B,C,D) registers receive the skew values from the SCNTR 518. After all the frames of a frame group have been received, the SKEW 501 (A,B,C,D) registers contain a measurement of the relative arrival times of the individual frames of the frame group. The SKEW 501(A,B,C,D) registers that contain a zero value indicate which frames arrived first.

The SKEW 594 register is not decremented by the microprocessor as the skew measurement procedure progresses, but is kept at a value of 255 throughout the entire procedure. Each time a frame group is received, the LC and summary information are stored in the FIFO 568. If some of the frames of the frame group are damaged by noise, they are not recognized as frames and the SCNTR 518 eventually reaches a value of 255 (equaling the value in the SKEW 594 register). The LC and summary information is then written into the FIFO

568. Writing into the FIFO 568 causes an interrupt to the microprocessor. The microprocessor checks the summary information to see if all the frames of the frame group have been received. If all of the frames have been received, the values in the SKEW 501(A,B,C,D) registers is saved in the microprocessor's local store before a validity test of skew data is made.

FIG. 11 shows the Skew Table 1102 which is kept in the microprocessor's local storage, and there is an entry for each of the four receivers. For each receiver, there is a minimum (MIN) 1104 and maximum (MAX) 1106 variable, and each of these variables is an 8 bit number having possible values from 0 to 255. As skew measurement frames are received, the values in the Skew Table 1102 are processed by the microprocessor.

The skew measurement test proceeds in exactly the same way as shown in either FIG. 6 or the alternate method shown in FIG. 9. The same skew-measure-request and skew-measure-response frames are used as the calibration frames with the improved method; however, many fewer frame exchanges are needed to achieve the same accuracy as are required by the previous methods. The flowcharts shown in FIGS. 7, 8, and 10 are no longer used.

FIG. 12 is a flowchart of the improved skew measurement code based on the frame exchange shown in FIG. 9. The frame exchange shown in FIG. 6 would result in a procedure that is largely the same as the procedure described here. The same procedure is used by both ends of the link with the exception of the handling of the first skew measurement frames. These differences are described below. In the Initialization 1202 step at the beginning of the skew measurement procedure, all of the receivers are active, the SKEW 594 register is set to 255, and the Skew Table 1102 is initialized. The MIN 1104 variables are set to 255 and the MAX 1106 variables are set to zero. Initializing the MIN 1104 and MAX 1106 variables to these values may seem contradictory, but as demonstrated below, the intended use of the MIN 1104 variable is to determine the minimum skew value as the skew procedure proceeds and the intended use of the MAX 1105 variable is to determine the maximum skew value as the skew procedure proceeds. The MIN 1104 variable starts at the largest value and is reduced as the skew values are processed, and the MAX 1106 variable starts at the smallest value and is increased as the skew values are processed.

The Send Measure Frame Group 1204 step is executed from the Initialization 1202 step only if this side of the link has the higher ND. If this side of the link has the lower ND, it waits for a skew-measure-request 906 frame group from the higher ND side of the link. In the Frame Received 1206 step, the code is waiting for an interrupt caused by the receipt of a frame group. After a frame group is received (either a skew-measure-request 906 or skew-measure-response 908), the code checks to see if all of the frame group has been received in the Frame Complete 1208 step. If all of the frame group has not been received, the microcode pauses in the Wait 1210 step for about 1 millisecond before it decides that the missing frames of the frame group will never arrive. This one millisecond wait time allows for a possible fiber optical cable length mismatch of up to 200 kilometers, enough time to account for any miscalculation. If a late frame is received within the one millisecond as determined in decision block 1209, the corresponding receiver is outgated in step 1211 and the mea-

surement procedure continues. If a late frame is never received, the microcode assumes that the missing frame was damaged by link noise, and either a frame-measure-request or frame-measure-response frame-group is sent to the other side of the link depending on the weather this side of the link has the higher or lower ND. By assuming that one or more of the frames was lost, the receiver is not outgated by a lost frame and is allowed to continue to participate in the skew measurement procedure. If the frame group is complete, the SKEW 501(A,B,C,D) registers are sampled in the Read Skew Registers 1212 step, and the sampled value is stored in the microprocessor's local storage.

The Skew Regs Still Valid 1214 step is a test to ensure that exactly one frame has been received from each receiver. If more than one frame is received, the value in the SKEW 501 (A,B,C,D) registers is meaningless since the additional frame changed to value in the SKEW register. In this embodiment the test consists of checking the V0 528, V1 530, V2 532, and V3 534 bits for a value of zero and making sure that the FIFO 568 is empty. If either of these conditions is not met, the other end of the link is not following the skew measurement protocol in that it sent an additional frame before allowed to by the protocol. The test is then restarted from the beginning.

The Update Table 1218 step is shown in more detail in FIG. 13A. For each receiver, the values obtained from the SKEW 501 (A,B,C,D) registers are compared to the MIN 1104 and MAX 1106 variables in the Skew Table 1102. First, if the skew value is less than the MIN 1104 variable (shown in step 1302), the contents of the MIN 1104 variable is replaced with the skew value (shown in step 1304). Otherwise, the MIN 1104 variable is left unchanged. Second, if the skew value is more than the MAX 1106 variable (shown in step 1306), the contents of the MAX 1106 variable is replaced with the skew value (shown in step 1308). Otherwise, the MAX 1106 variable is left unchanged. The above procedure repeats for the other configured receivers (as many as three) before the Update Skew Table 1218 step completes.

After the Update Skew Table 1218 step completes, a counter indicating the number of successful skew-measure-request or skew-measure-response frame-groups that have been received is incremented in the Increment Count 1220 step. This count is tested in the Done 1222 step, and control is passed to the Send Measure Frame Group 1204 step if more tests are required. The Complete 1224 step checks if the other side of the link has received enough skew measure frames. If it has not, control is passed to the Send Measure Frame Group 1204 step. If the other side of the link has completed, the microcode the entire skew test has completed.

The Process Table 1226 step is shown in more detail in FIG. 13B. First, all of the variables are converted from cycle counts to link word times in the Convert 1350 step. Next, in decision block 1352, for each receiver the MIN 1104 variable is subtracted from the MAX 1106 variable. If the difference is greater than twice the random skew value (RSKEW) allowed by the system, an error 1354 is signaled. In this case, the link is considered unreliable, and operations cease. In actual practice, this difference is computed during testing in engineering test of the system in a laboratory environment when the computation is made many many (e.g. a million) times. If the system is designed correctly, this difference is always within acceptable bounds, and the

MIN 1104 variables are added to twice the random skew (RSKEW) value in step 1356 to determine the value for the maximum skew. This method is used to determine the maximum skew because even with multiple frame-group exchanges, there is a finite probability that the maximum skew value may not occur during the test. With this method, multiple frame-group exchanges are used to find the smallest minimum skew so that the required frame to frame times are minimized. After the maximum values are computed from the MIN 1104 variables, they are checked for values greater than 127 word times in decision block 1358. If any are larger than 127, their corresponding transceivers cannot be used to form an operational link. In this case, the SKEW value is replaced by a value of 255 in step 1360.

Returning to FIG. 12, the Send Results Frame 1228 step sends the values computed in the Process Table step 1226 in the skew-results-request 922 and skew-results-response 924 frames. Recall that any values that are larger than 127 are replaced with a value of 255 indicating that the corresponding receiver was either outgated or had excessive skew.

After the exchange of the skew results frames, the establishment of the link proceeds as described earlier.

While we have described our preferred embodiments of our invention, it will be understood that those skilled in the art, both now and in the future, may make various improvements and enhancements which fall within the scope of the claims which follow. These claims should be construed to maintain the proper protection for the invention first disclosed.

What is claimed is:

1. A method for measuring and verifying the skew between frames of a frame group over a communications link having a plurality of carriers, comprising the steps of:

receiving, in parallel, a series of calibration frames over each of the carriers;  
measuring the relative arrival times of said calibration frames;  
constructing a skew table having measured skew for each carrier;  
updating said skew table as successive measurements are made by receiving successive calibration frames;  
computing the maximum skew from said skew table; and  
communicating said computed maximum skew to the opposite end of the link.

2. The method of claim 1 wherein said measuring comprises recording in a group of registers the relative arrival times of the calibration frames as they are received.

3. The method of claim 2 wherein each of said registers is set to the value of a corresponding skew counting register as said calibration frames are received.

4. The method of claim 1 wherein the skew table is constructed at the beginning of the skew measurement procedure and has one set of entries per receiver.

5. The method of claim 4 wherein said skew table entries are initialized with a maximum value in the variable that records the minimum value for the measured skew and a minimum value in the variable that records the maximum value for the measured skew.

6. The method of claim 1 wherein said skew table is updated as successive calibration frames are received such that a comparison is made between the measured

21

value and the corresponding value in the skew table for each of said carriers further comprising the steps of:  
 comparing the measured skew value to the corresponding skew table minimum value;  
 replacing the skew table minimum value if said comparing reveals that the measured skew value is less;  
 comparing the measured skew value to the skew table maximum value; and  
 replacing the skew table maximum value if said comparing reveals that the measured skew value is more.

7. The method of claim 1 wherein the final skew values are computed from the skew table at the end of the skew measurement procedure further comprising the steps of:

adding twice the value of the random element contribution to the skew to the minimum numbers in the skew table; and

replacing values over 127 with 255 to indicate that the receiver is outside of the allowable skew limit.

8. The method of claim 1 wherein in a manufacturing process for computing a maximum random skew, all cycle count variables for measuring skew are converted from cycle counts to link word times,

and, in a decision block for each receiver a MIN variable is subtracted from a MAX variable, and if the difference is greater than twice the random skew value (RSKEW) allowed by the system, an error is signaled, in which case, the link is considered unreliable, and operations cease.

9. A transceiver for sending and receiving message frames, comprising:

means for coupling said transceiver to a communications link;

transmission section for asynchronously transmitting message data in frame groups across multiple car-

22

rier conductors of the link and for coupling the frame groups together;

a receiver section for receiving frames and for determining which frames are part of a group and for determining a relative skew on a conductor of a link by performing skew measurements on a series of calibration frames sent by a transmitter on another end of said link;

a monitor means for performing said skew measurements, including a plurality of skew measurement registers for measuring the relative skew between frames of a frame group, each register having an input from a skew counter, and outputs coupled to a microprocessor controlling the monitor means, each register gating in the value in said skew counter at the time a link control word is received; said monitoring means including means for gathering, each time a link control word is received for a skew measure calibration frame, skew information from all of the carrier conductors of the link, means for measuring the relative arrival times of said calibration frames;

means for constructing a skew table;

means for updating said skew table as successive measurements are made by receiving successive calibration frames; and

means for computing the maximum skew from the said skew table and for communicating said computed maximum skew to the opposite end of the link.

10. A transceiver according to claim 9 further comprising:

means for determining that one or more frames of the calibration frames have been lost and have not exceeded the maximum measurable skew limits.

\* \* \* \* \*

40

45

50

55

60

65